

Exploring Little's Law

Class handout

[As class starts, Radu passes out a handout with the following case and asks the students to read it out loud as a group, each taking turns to read. Radu gives students some time to ponder and discuss among themselves in smaller groups.]

Case: Changing Lanes at Metrovia Central

Last year, the Metrovia Transit Authority (MTA) hit a milestone: over 5 million daily passengers. Much of the city's rapid growth was absorbed by the metro network, and it was operating at near-maximum capacity during peak hours. Despite upgraded trains and expanded platforms, one recurring chokepoint remained: vertical movement between platforms and street level.

At Metrovia Central, the city's busiest hub, long escalators carried passengers up from the deep train platforms to the main concourse. Three adjacent escalators served this upward flow, and MTA policy—like in many cities—was to encourage passengers to stand on the right and walk on the left.

But during the 8:00–9:00 a.m. rush, congestion on the platform was crushing. Crowds built up at the bottom of the escalators, delaying platform clearances and slowing boarding on incoming trains. A station supervisor noted, “It’s not the trains that are the problem. It’s these blessed escalators.”

An MTA capacity analyst, Sofia Chen, had recently returned from a research visit to East Asia, where she observed escalator etiquette in several cities. In some of these systems, riders stood calmly on both sides. Intrigued, Chen conducted a simulation based on Metrovia Central's escalator specs. Her model suggested that changing the protocol to “stand on both sides” could increase throughput during peak periods.

The MTA commissioned a three-week pilot on two of the three escalators. The process ran as follows: Riders exited trains and moved toward the escalator bank. Staff then guided them to the next available escalator. Escalators A and B had signage and audio announcements: “Please stand on both sides.” Riders on Escalator C followed the traditional stand-right/walk-left convention. Station staff observed flow rates, rider behavior, and queuing times.

Operationally, the results were promising. Escalators A and B handled more passengers per minute than Escalator C, and station managers noted fewer blockages on the platform. But passenger reactions were mixed. Some complied without comment. Others voiced frustration. A few attempted to walk anyway, weaving between standers or nudging past. Overall, staff felt the trial was a success, but were not sure about rolling the policy out across the whole network.

Question for consideration:

Should the MTA expand the standing-only configuration to other stations?

Class discussion

RADU: Before we start, Ming, you look frustrated again. Is this going to happen every class?

MING: Well, that's up to you, isn't it?

RADU: Apparently so. So what is bothering you today?

MING: I don't see how we can even hope to answer this question.

RADU: Why not?

MING: My first mentor had a saying, "In God we trust, all else bring data". There is no data in this case about these other stations; how can we possibly decide what the MTA should do?

RADU: That is not a quote of your mentor; that is an old saying usually attributed to W. Edwards Deming, and even that is possibly misquoted. But let's cross that bridge another time. Ming, we are not actually making this decision for real! Of course, we would collect more data if this were the real situation. In class, we are simply wondering what we would need to consider to answer this question, examine the tradeoffs, and argue with one another a bit. Is that okay for you?

MING: But that is all shallow and perfectly pointless if not based on real data. Then we really are just in a classroom arguing and pretending we are important.

RADU: Who said we are important? Ming, we are learning. We can learn via thought experiments. Einstein did it and so can we. *[Pauses.]* But I know where you are coming from. How can we make this more real? Maybe I can ask you, what data would you want to collect to feel comfortable with providing your opinion?

MING: So we are just going to make up data now? That's even worse than having no data.

RADU: I am not asking you to make up data; I just want to know the type of data you need. *[Stops, thinking what to say next.]* Did you ever see an equation like this before? *[Writes on the board.]*

$$y = mx + b$$

MING: Sure, that's an equation of a line.

RADU: What is the point of it? There is no data in it, so how could it be useful?

MING: Well, it's a formula, which tells us what data we might need. For example, we can think of m and b as representing a place to insert data.

RADU: I like that you said this is a formula. What is this formula trying to say?

MING: It means there is a linear relationship between x and y , and the coefficient m and b specifies this relationship.

RADU: True, true. So you see, here, what matters is this linear relation. In this formula, even though I don't give you numbers, can you still analyze the behavior of this line. Ming?

MING: What do you mean "analyzing the behavior of this line"? That sounds a little fancy, doesn't it? I just remember m controls the direction of the line, like whether it goes upwards or downwards, and b controls the location of the line.

RADU: What you just said is what I mean by analyzing the line. That makes you "a little fancy" indeed. *[Everyone else class laughs. somewhat uncomfortably.]* So Ming, do you see this? I didn't give you any number, but you still manage to analyze something about this formula. This type of analysis is actually more general because it is not tied to any given numbers.

MING: Okay, okay, you made your point. Let's move on.

RADU: If you say so. *[Everyone laughs]* So we are trying to figure out whether the standing-only configuration is good in that situation. It seems to be good in the station they tried, but what about other stations? Is there a " $y = mx + b$ " type formula that we can reason about here that can help us?

FARAH: Yes, Little's Law!

RADU: Very good! After all, we only have a couple of formulas that we have discussed. Where are you seeing Little's law in here, Farah?

FARAH: Ultimately, the issue MTA is facing is that the inventory of travelers on the platform is too large. And the reason it is too large is that the escalators are too slow. What we really need to ask is: does the standing-only policy make the platform empty out faster by speeding up the escalators? We can use Little's Law to make sense of this.

RADU: Do you agree Binesi, that the question boils down to how fast the escalators are?

BINESI: Hmm, not fully. What do you mean by "fast"? I think we are not proposing to change the speed of the escalators, like how fast the stairs move, although come to think of it, that might be a good idea.

RADU: Very good point, we need to be a little careful with our language here. *[Pauses.]* But making the escalators move faster would indeed help, wouldn't it? I never thought of that before. *[Looks towards the ceiling, thinking.]*

BINESI: Yes, because we could move more people off the platform that way.

RADU: Certainly. Is that what you meant by "speeding up the escalators", Farah?

FARAH: Oh, I used "speed" colloquially; I meant that we want to see if the standing-only policy increases the *throughput rate* of the escalators, not by changing the physical speed of the escalators, but by adjusting how many people are on the escalators. The case said "standing only" is faster at moving people in this one station, so I think we should understand why that is.

BINESI: Agreed, and I would assume the escalators are going as fast as they can really, keeping safety considerations in mind, of course.

RADU: Yes, let's assume that. So I see now Farah, how you were thinking about Little's Law: you have some feeling about " R " in the $I = RT$ formula. In a way, we are asking how changing the policy for standing-only or walking changes " R ", is that right?

FARAH: Exactly.

RADU: Can we call them R_W for throughput of a walking policy and R_S for the throughput of standing-only?


[Pauses for a moment.]

Can we step back a bit before we start arguing which of R_W and R_S is bigger? We dove into this Little's Law perspective pretty quickly, and I want to make sure we have not lost the forest for the trees. Haoyu, how would a process flow diagram look in this scenario here?

HAOYU: I don't ride the subway.

RADU: Of course not. *[Everyone laughs.]* But I think you can picture it. Please, draw a process flow diagram for what is happening here.

[Haoyu gets up and draws on the board]



[Everyone claps when he is done.]

RADU: Great job Haoyu. I like your drawings of the train and escalators. This is a nice visualization of what is happening, but I am not sure if it is what I would call “process flow” diagram. I see your inputs are passengers coming off the train and the outputs are them leaving at the top of the escalator. But what are the resources and activities here? Where are your boxes and arrows?

HAOYU: Honestly, I am a little confused about that. Because the activities seem to be mostly waiting around, either on the platform or on the escalator itself. So I just decided to draw it.


RADU: We like your drawing! Very inspiring. But yeah, there feels something kind of “degenerate” about this situation. Before we had some worker doing something to help a customer, the worker was a resource and what they did to serve was the activity. Can you try to apply that perspective here?

HAOYU: Well, I guess the closest thing to a “worker” is the escalator. And the action is “moving”?

RADU: Precisely! Now, why did you doubt yourself? Here, can I draw something, Haoyu?

HAOYU: Sure, go ahead. *[Haoyu sits down as Radu moves to the board.]*

RADU: Are you saying something like this?




HAOYU: Yeah, that looks right.

RADU: Can someone help me flesh out the diagram more?

CAMILA: Your figure is missing the pivotal concern, the platform! I want to suggest where to put that, but I am not sure myself. Is the platform a resource?

RADU: What do you think?

CAMILA: Well, I know there is a waiting activity, and every activity needs a resource to do it maybe?
Is the platform the resource doing this activity?
[Radu draws while Camila speaks]



RADU: Yes, that is a nice way to think about it. There is this platform, and it is a pool of “spots” for people to stand while waiting, and if there is no “spot” for you, you somehow cannot wait, you can’t even get off the train!

[Pauses for a moment looking at the figure on the board.]


Now, when I look at this, I only see the platform and escalators. But the intervention we are looking at has to do with the standing-only policy. But how does the policy directly change the platform or the escalator? How could changing the policy change this figure at all?

MING: You don’t have to school me twice on the same point, Radu. It’s just like a few lectures back. We are missing the people; the people are also a resource.

RADU: How so?

MING: The people are a resource in that “Ascend” action. The people combine with the escalator to “Ascend”. Nothing meaningful ascends if there are no people on the escalator.

RADU: Well put. [Adjusts the picture.]



I guess this is a key observation behind pondering whether R_S or R_W is greater. These are only different if people’s behavior changes the throughput, right Ming?

MING: That’s my point.

RADU: Indeed it is.

MING: But now don’t we need to collect data to decide which one is bigger?

RADU: [Smiles] Again with the data, Ming? Of course, you are right in a certain sense, but let’s reason instead with a formula. Can we get a formula for R_S and R_W ?

FARAH: Yes, we certainly can! From Little’s Law, we have: $R = \frac{I}{T}$. So we should have $R_S = \frac{I_S}{T_S}$ and $R_W = \frac{I_W}{T_W}$.

[Radu draws the formulas on the board.]

$$R_S = \frac{I_S}{T_S}, \text{ and } R_W = \frac{I_W}{T_W}$$

RADU: I love that. You just took Little's Law, " $I = RT$ ", and divided both sides by T to get $R = I/T$. The Little's Law formula opens like a window into our exploration of R_S and R_W . Binesi, how does this formula help you think about this question of whether R_S is bigger than R_W , or vice versa?

BINESI: Well, this really depends on people's behavior, doesn't it?

RADU: As always. Continue.

BINESI: I want to think about the T 's first; that seems easier [*Points to the bottom of the two formulas on the board.*] I can't help but think that T_S is bigger than T_W . If you are standing, it takes longer to get up the escalator. That's the whole point.

RADU: I am with you on that; it seems like necessarily true that $T_S \geq T_W$ [*Adds this to the board.*]

$$R_S = \frac{I_S}{T_S}, \text{ and } R_W = \frac{I_W}{T_W} \text{ with } T_S \geq T_W$$

Do we need data for that Ming?

MING: I would have to admit that I would believe it without data.

RADU: So going back to the Deming quote, Binesi is a "God"? You trust her without data? [*The class laughs and students look at Ming.*]

MING: I just mean that it is logical.

RADU: Ah yes, logic! There is something called logic, and logic is different from data, what do you think Ming?

MING: [*Impatiently.*] Let's just keep going.

RADU: Let's. Can anyone use their "logic" to tell me about I_S and I_W ?

[*A brief pause as students think about the question.*]

BINESI: When people are standing, you can "pack" them into the escalator better. When you are standing straight, you can fit more people around you comfortably. But when you are walking, people need to steer clear a bit; otherwise you might kick them or they might trip you. So I think $I_W < I_S$.

RADU: Very good. How do you know this Binesi? Is that based on data? Or logic?

BINESI: I guess it's my own experience. If I am walking behind someone, I stay farther away from them. I have noticed others do that. I am not sure whether to call that data, logic, or just common sense.

RADU: What do you think Ming? You are the one with a keen eye for all things "data".

MING: I have a similar experience to Binesi. So maybe a couple of data points, her experience and my experience, can constitute "data".

RADU: I'd buy that. So let's write it down.

$$R_S = \frac{I_S}{T_S}, \text{ and } R_W = \frac{I_W}{T_W} \text{ with } T_S > T_W \text{ and } I_S > I_W$$

RADU: But now we have come to a crossroads with our formulas, haven't we Farah? Can "logic" or "experience" alone tell us which one of R_S or R_W is bigger?

FARAH: Sadly, no. The numerator of R_S in the formula is larger than R_W 's numerator, but the denominator is also bigger. This makes the overall size of R_S in comparison to R_W uncertain.

[Radu writes on the board to illustrate this.]


$$R_S = \frac{I_S}{T_S}, \text{ and } R_W = \frac{I_W}{T_W} \text{ so } R_W \stackrel{?}{>} R_S$$

RADU: Precisely. We need to think a little more. Let me ask you Camila. If we had a subway station with a really long escalator. The station is way down there. Would that change things?

CAMILA: Hmm, I feel it does. If the escalator ride is really long, I think fewer people will choose to walk. This means that I_W might be smaller than we thought because it is not popular, so few people will choose to walk. In this case, I feel that R_S is more likely to be bigger than R_W .

$$R_S = \frac{I_S}{T_S}, \text{ and } R_W \downarrow = \frac{I_W \downarrow}{T_W} \text{ with } T_S \geq T_W, \text{ and } R_S \geq R_W$$


RADU: Right, I really like how you highlighted that there is a choice being made by travelers as they approach the escalators. Maybe we should adjust the picture to reflect that. [Erases and redraws the process flow figure.]



To Camila's point, the percentage of people who make the "walk" decision impacts the throughput of walk escalators.

FARAH: I also think it could also depend on the time of day.
 RADU: How so?
 FARAH: In the morning, people are late for work, and more might choose to walk to make up time, rather than stand.
 RADU: Oh interesting. I hadn't thought of that ...
 BINESI: For me, I would be more urgent to leave work to get home to my family.
 RADU: Everyone has their priorities, indeed.
 HAOYU: Actually, I have a question. I kind of don't get why you wrote the process flow diagram that way. I have my own; do you mind I show it?
 RADU: Of course not, please draw it for us.

[Haoyu gets up and writes on the board.]



HAOYU: I think that you decide to walk or not before you join the line, no? So shouldn't it look like this?
 RADU: Hmmmm, I see ...
 BINESI: Does this make any difference? I am not seeing why this is a big deal.
 HAOYU: Yeah, I am not sure if it does. I just see it this way. I decide whether to walk before I join the line.
 RADU: Who can help us distinguish these two?
 MING: I see them as very different, and I believe the single queue one is more appropriate in the situation where the station is really busy. When you get out of the train, if the platform is really busy, you will be waiting in line right away to exit. There is just one big line there. I can change my mind whether to stand or walk closer to the end of the line.
 HAOYU: Oh, I see your point. My picture would make more sense if I had some choice of which line to join, which suggests the platform is not super busy and I can distinguish the two lines from the beginning. But you are right, this is not easy to do in truly crowded situations. I often have no idea where the lines start and for which escalator when I first exit the train in a really busy station.

RADU: This all makes sense, and maybe there is some dispute and interpretation. But can we stick with the first picture for the sake of argument, the one with a triangle before the diamond? [Everyone nods.] Good, so let me ask you all this. Is it possible for someone running late to prefer a standing-only policy for all escalators?

CAMILA: If I am late, I will always walk, I will never prefer a standing-only policy. Stopping me from walking is impeding my freedom to choose, and that's not right.


RADU: Your freedom to choose here is more valuable than being late to work?

CAMILA: But choosing to walk makes things faster, so aren't I less likely to be late if I walk?

RADU: How can you be sure of that? Suppose we are in a situation where $R_S > R_W$, like maybe we are in a situation where the escalator is really long. Now, you get off the train at a really busy time; you can barely get on the platform. Can you start walking on the escalator right away?

CAMILA: No, I need to wait until I get to the escalator.

RADU: Ah yes, you need to wait. You are here in the picture. [Draws a stick figure representing Camila.]



Given you are here in the triangle, would you prefer the people ahead of you in line to walk or stand on the escalator?

CAMILA: Well obviously ...no wait ...[Looks puzzled.]

RADU: [Smiling] I warned you about saying “obviously”, didn't I?

CAMILA: Don't rub it in. [Everyone laughs.]

RADU: So would you prefer others to walk or stand?

CAMILA: I am going to say “stand”.

RADU: And why?

CAMILA: Well, as you said, R_S is bigger than R_W . If that's the case, the platform will empty out faster in front of me, and I can get to the escalator sooner.

RADU: This is very astute Camila. You used another formula we have talked about before. [Turns to write on the board.]

$$T = T_{\text{process}} + T_{\text{wait}}$$

Here “process” is the “Ascend” activity. But waiting is the bigger proportion of time here, is it not? Walking makes T_{process} faster, but others walking makes T_{wait} slower. Does everyone see that?

BINESI: Can you say it again?

RADU: Camila, maybe you can speak to that. I said it in my way, but you were the one who said you preferred others to stand. Can you explain again why?

CAMILA: Here is how I thought about it. I want to get to the front of the line as soon as possible, yes? I can use Little’s Law to determine my wait. My time in the line is the number of people ahead of me in line divided by the throughput of the line.

RADU: *[Interrupting]* Can I stop you there, Camila, something like this? *[Writes on board.]*

$$T_{\text{time to front of line}} = \frac{I_{\text{ahead in line}}}{R}$$

CAMILA: Yes, exactly. Now, the number of people ahead of me in line, which is this $I_{\text{ahead in line}}$ thing, doesn’t change with the policy, but this R does. I want R to be larger to shorten my wait. So I prefer it if others stand since $R_S > R_W$.

RADU: I like this, yet another use of Little’s Law! That is what I was telling all of you; it shows up everywhere. Binesi, does Camila’s explanation clear this up for you?

BINESI: Yes, it does. That makes a lot more sense. But doesn’t this create a kind of conflict in you Camila. You want others to stand, but when you get to the escalator, you walk. Wouldn’t the same logic apply to everyone?

CAMILA: Well ...*[Pauses.]* I see your point, everyone running late will want to walk once they get to the front of the line.

RADU: And now do we all see why this standing-only policy is hard to enforce?

MING: I do. People are selfish; this will never work. Everyone will just walk.

RADU: But if everyone walks, everyone is worse off. People want this?

MING: People just want others to be worse off. *[Everyone laughs.]*

RADU: But we ourselves get trapped in it. If others are worse off, we are necessarily also worse off. We are all stuck in the same triangle on the platform.

MING: But at least no one is happy.

[Everyone laughs.]

RADU: On that positive note, we should wrap it up here.

Email

Dear students, a lively class today. I want to use this email to develop a few threads that came up in the discussion more carefully.

Reasoning without data: the oldest trick in science. I want to return to the sparring match Ming and I had at the start of class, because there is something important underneath it. Ming wanted data before she would believe anything, and she is right that data matters enormously. We will spend much

of this course respecting it. But notice what we actually did in class: we reached pretty firm conclusions about the two escalator policies without a single measurement of the real platform. We reasoned from the structure of Little's Law alone and our own experiences.

We tend to think of formulas as machines for producing answers. You put numbers in, you get a number out. Ming's frustration today was, in part, a frustration of this type: without specific data for R_S , R_W , I_S , I_W , T_S , T_W , what can we possibly do with the formula $I = RT$?

Quite a lot, it turns out. The power of a formula like Little's Law is not only that it produces answers. It is that it structures thinking and builds relationships between different quantities. Consider the analogy from class: the equation $y = mx + b$ contains no data whatsoever. And yet simply by thinking about the role of m and b —what happens when m is large or small, positive or negative; what b does to the position of the line—you can understand the behavior of an entire family of lines, without a single data point in hand. Little's Law works the same way. It tells you which quantities matter, how they relate to each other, and therefore what tradeoffs you should be reasoning about. Data is needed to answer specific numerical questions. Logic and formulas are what you use to identify which questions are worth asking, and to determine whether a data-free argument can settle the matter.

This is what I mean when I say Little's Law is a window into the problem. It is not the answer. It is a structure for finding where the answer can come from and tells us how the change in one quantity affects another.

Reasoning by logic before observing data is not a trick I invented for a business-school classroom. It is one of the founding moves of modern science. In 1638, Galileo staged an argument against a belief that had stood since Aristotle: that heavier bodies fall faster than lighter ones [3]. He did not climb a tower and drop anything. He simply asked his readers to imagine tying a heavy stone to a light one. If the light stone falls slower, it should drag on the heavy one and slow the pair down—so the pair falls slower than the heavy stone alone. Yet the pair, taken together, is heavier than the heavy stone alone, so by Aristotle's own rule it should fall *faster*. The same object, falling both slower and faster than itself: a contradiction, reached with no apparatus, no stopwatch, no data at all. The belief simply could not be consistent, and that was the end of it.

God and data. Ming rallied to a famous slogan: “In God we trust; all others must bring data”. I flagged in class that it is usually hung on W. Edwards Deming, and possibly misquoted. Let me close that loop, because the truth is more interesting than the slogan. The line appears on mugs and posters of Deming everywhere, yet no one has ever produced a place where he actually said or wrote it. Its provenance is a small mystery, which is a delicious irony for a motto demanding that all claims come with evidence. The slogan itself cannot bring data as to its own origin.

More to the point, the documented Deming was far more cautious about data-worship than his bumper sticker suggests. In his own book he warned that “the most important figures needed for management of any organization are unknown and unknowable” [2], by which he meant that the quantities that matter most, such as the compounding effect of a single delighted or disgusted customer, often cannot be measured at all, and that a manager who attends only to what is easily counted is easily lead astray.

When logic runs out. Having now praised reasoning-without-data twice, I must guard against overselling it, because the opposite error is just as real. Some questions genuinely turn on a magnitude on how big a number is, and there no amount of elegant reasoning will save you from having to go and mea-

sure. The trick is developing a feel for which kind of question you are holding. Does the answer depend on the *shape* of a relationship, which logic can often supply, or on the *size* of a particular quantity, which only data can? Our escalator, as it happens, is a beautiful case where the answer flips depending on the situation. A large part of judgment in this field is knowing, for the problem in front of you, when logic stops, and measurement must begin.

Local versus global optimization, and the tragedy of the commons. One of the more interesting moments in class came when Camila realized she wanted everyone *else* to stand on the escalator, even as she planned to walk herself once she reached the front of the line. This is not a contradiction. It is entirely rational. But it is rational only from Camila's perspective in isolation. It falls apart if we ask what happens when everyone reasons the same way.

There is a name for this kind of situation. It is called the *tragedy of the commons*, a concept introduced by the ecologist Garrett Hardin in a 1968 paper of the same name. Hardin's original example was a shared pasture: each individual herder, acting in their own interest, has an incentive to graze one more animal, since they capture the full benefit of the extra animal while the cost of degrading the pasture is shared by all. Each herder's local reasoning leads them to graze more. But if all herders follow this logic, the pasture is overgrazed and destroyed, which is an outcome that is bad for every individual herder, including the ones whose local reasoning led them there.

The escalator is a nearly perfect small-scale analogue. Each commuter, once they reach the front of the line, has a private incentive to walk: it reduces their personal processing time T_{process} . But when everyone walks, the throughput R of each escalator falls, the platform inventory I grows, and the waiting time T_{wait} for everyone, including the walker. If the increase in T_{wait} is large enough to outweigh the decrease in T_{process} , the walker ends up worse off than if they had stood. And yet the rational choice, from the individual's vantage point just before stepping onto the escalator, is still to walk.

This gap between what is good for the individual (*local* optimization) and what is good for the system as a whole (*global* optimization) is one of the central tensions in operations management. It appears in supply chains, in hospitals, in traffic networks, in call centers, in almost any setting where individuals make decisions that interact with each other through a shared resource. We will return to it many times.

Notice, finally, how Little's Law helps clarify what is going on. Walking decreases T_{process} , the smaller part of total flow time. But it also decreases R , which, by $T_{\text{wait}} = I_{\text{ahead}}/R$, increases T_{wait} , the larger part. The phenomenon has a name: "hurry up and wait." We rush through the part of the process we control (ascending the escalator) while lengthening the part we do not control (standing on the platform). Operations management asks us to resist the instinct to optimize the visible, controllable piece of the process and instead ask what is happening to the whole.

Braess's paradox, or when taking away a choice helps everyone. Camila objected, reasonably, that a standing-only rule impedes her freedom to choose. I want to hand her a genuinely unsettling result that puts that freedom in a new light. In 1968 the mathematician Dietrich Braess showed that in a road network, *adding* a road, i.e., giving every driver one more option, can make everyone's journey longer [1]. Each driver rushes for the tempting new shortcut; the shortcut clogs; and the equilibrium in which everyone has chosen selfishly is worse, for every single driver, than the world without the extra road. The converse is the striking part: sometimes *removing* an option makes everyone better off. This is not merely theoretical, e.g., when New York City closed 42nd Street to traffic for Earth Day in 1990, planners braced

for gridlock and instead watched the flow improve.

The standing-only escalator is a small Braess move. Take away the option to walk on a tall escalator, and the whole system moves more people, including the very people who would have chosen to walk. So Camila's freedom is real, but it is the peculiar kind of freedom whose *universal* exercise leaves everyone worse off, and there the kindest policy may be to remove it.

What actually happened. Ming will be pleased to hear that our fictional case is not so fictional after all. In 2015 and 2016, the London Underground ran very nearly Sofia Chen's experiment for real, at Holborn station: for a stretch, they asked everyone to stand on both sides of certain "up" escalators, with no walking allowed. The result vindicated Sofia's simulation. Throughput rose by roughly thirty percent: one escalator that had carried about 2,500 people during the morning peak carried around 3,250 once it went standing-only. The station, which had been forced to close its gates against overcrowding several times a week, had to do so only once during the trial [4].

This worked at Holborn only because Holborn's escalators are unusually tall—about 23 metres—so tall that hardly anyone bothers to walk them, which leaves the whole walking side sitting empty and wasted. On a *short* escalator, where plenty of people do walk, forcing everyone to stand actually *reduces* capacity, because you have idled a lane that was doing real work; a trial on the much shorter escalators at Canary Wharf found precisely this. Whether banning walking helps or hurts depends on the height of the escalator. Logic tells you the effect *can* go either way; only the number tells you which. The London Underground did not, in the end, keep the policy—because sustaining it required continuous staff coaxing that proved impractical. The class worried the policy would be hard to enforce. The class was right.


A worked example: the laundromat. Let me work through a more involved application of Little's Law, in a setting where two subsystems are nested inside one another, which is how things usually look in practice.

Example: Suds & Spin Laundromat

Suds & Spin is a busy self-service laundromat. Customers arrive at an average rate of 30 per hour. Some customers glance through the window, see that machines are nearly all occupied, and leave without entering, which we call *balkers*. On average, 12 customers are waiting inside for a machine to become free, and the average wait for a machine is 30 minutes. There are 8 washing machines, and on average they are 75% utilized. All customers who enter eventually get a machine.

1. At what rate do customers balk?
2. On average, how long does a customer spend using a machine?

Here is a process flow diagram for this situation. The dashed arrow represents the balking path—customers who leave without entering the system.



We can think of the laundromat as two subsystems in series, each with its own inventory, throughput, and flow time. Finding the boundary of each—what counts as “arriving” and “leaving”—is the key first step.

System 1: the queue. The flow unit is a customer waiting for a machine. We are given:

- Average inventory in the queue: $I_{\text{queue}} = 12$ customers.
- Average time waiting in the queue: $T_{\text{queue}} = 30 \text{ minutes} = \frac{1}{2} \text{ hour}$.

Applying Little’s Law:

$$I_{\text{queue}} = R_{\text{enter}} \cdot T_{\text{queue}} \quad \Rightarrow \quad R_{\text{enter}} = \frac{I_{\text{queue}}}{T_{\text{queue}}} = \frac{12 \text{ customers}}{\frac{1}{2} \text{ hour}} = 24 \text{ customers per hour}.$$

So 24 customers per hour enter the queue.

Question 1: the balking rate. Of the 30 customers arriving per hour, only 24 enter. The rest balk:

$$30 - 24 = 6 \text{ customers per hour}.$$

System 2: the machines. Since all customers who enter the queue eventually get a machine, the throughput into the machine system equals the throughput out of the queue: $R_{\text{machines}} = 24$ customers per hour.

The 8 machines are 75% utilized on average, meaning:

$$I_{\text{machines}} = 8 \times 0.75 = 6 \text{ customers}.$$

Applying Little’s Law:

$$I_{\text{machines}} = R_{\text{machines}} \cdot T_{\text{machines}} \quad \Rightarrow \quad T_{\text{machines}} = \frac{I_{\text{machines}}}{R_{\text{machines}}} = \frac{6 \text{ customers}}{24 \text{ customers/hour}} = \frac{1}{4} \text{ hour} = 15 \text{ minutes}.$$

Two things are worth pausing on. First, notice that the output rate of the queue became the input rate of the machine system. This chaining of subsystems is how Little’s Law handles more complex operations. Second, the total time a customer spends in the laundromat (if they don’t balk) is $30 + 15 = 45$ minutes, of which two-thirds is waiting and only one-third is actually doing laundry. A familiar feeling.

References

- [1] Dietrich Braess, Anna Nagurney, and Tina Wakolbinger. On a paradox of traffic planning. *Transportation Science*, 39(4):446–450, 2005. English translation of Braess’s 1968 paper.
- [2] W. Edwards Deming. *Out of the Crisis*. MIT Center for Advanced Engineering Study, Cambridge, MA, 1986.
- [3] Galileo Galilei. *Dialogues Concerning Two New Sciences*. Elzevir, Leiden, 1638. Translated by Henry Crew and Alfonso de Salvio, Macmillan, 1914.
- [4] Transport for London. Report on holborn pilot for standing on both sides of escalators. London Underground, 2016. Standing-only escalator trials at Holborn station, 2015–2016.

Extensions

[Posted as homework on the course’s website for this class session.]

1. Suppose it costs you c dollars to procure a pineapple and you can sell each pineapple for p dollars, where $p > c$. Demand is so high that you can sell every pineapple you procure at the full price p . How many pineapples should you procure? Do you need specific values of p and c to answer this?
2. Now suppose that selling a pineapple earns you p dollars only if it sells within two days of purchase, and otherwise earns you nothing (the pineapple spoils). Does the answer change? What additional information would you now need to make the decision? What does this exercise suggest about when formulas and logic suffice, and when numbers become necessary?
3. Recall during the class discussion, that Haoyu drew a second process flow diagram, which the rest of the discussion ignored. If the overall operations looked more like Haoyu’s process flow diagram, would the conclusions of the discussion be much different?
4. The blanket standing-only rule was good for the crowd but hard to enforce and locally resented. Design an alternative that captures some of the throughput gain while ducking those two problems.
 - (a) Specify your policy precisely enough that a station manager could put it in place tomorrow.
 - (b) Trace it through Little’s Law: which of I , R , and T does it move, and for whom?
 - (c) Stress-test your own design. Describe the first way a self-interested rusher would defeat or game it, what you would have to add to close that hole, and whether closing it quietly reintroduces the very enforcement problem you were trying to escape.
5. A popular coffee shop on campus serves customers continuously from 8:00 am to 6:00 pm. During the morning rush (8:00–10:00 am), customers arrive at a rate of 60 per hour, and the average number of customers inside the shop at any moment is 20. During the rest of the day, the shop is quieter: customers arrive at 20 per hour, and the average number inside drops to 10.

- (a) What is the average flow time (from entering the shop to leaving with a drink) during the morning rush? During the rest of the day?
 - (b) What is the total number of customers served over the full 10-hour day?
 - (c) If the shop manager wants to reduce the average flow time during the morning rush to under 15 minutes without changing arrival rates, what would need to change, in terms of the components of Little's Law?
6. A consultant tells a hospital CEO: "We analyzed your emergency department and found that your average patient wait time is 3 hours. If you hire two more triage nurses, you can cut that to 90 minutes." The CEO asks you to evaluate this claim before signing the contract. Using Little's Law, what specific data would you demand from the consultant before accepting or rejecting the claim?
7. Your company runs a call center. Average handle time (the time an agent spends on a call) is 6 minutes. You have 20 agents. On average, 8 callers are on hold waiting for an agent at any moment. Management wants to cut the average hold time in half.
- (a) Using Little's Law, determine the current throughput rate of the call center and the current average hold time.
 - (b) Management proposes two interventions: (i) train agents to cut handle time from 6 minutes to 4 minutes, or (ii) hire 10 additional agents. Use Little's Law to predict the effect of each intervention on average hold time, *assuming* the arrival rate of calls stays constant.
 - (c) Now suppose that cutting hold time causes more callers to *stay on the line* rather than abandoning the call—so the effective arrival rate into the system rises. Does this change your recommendation? What does this suggest about the limits of using Little's Law in isolation to predict the effect of an operational intervention?
8. A city's housing authority tracks units of affordable housing. At any given time, there are 2,400 households on the waiting list. Each year, 300 new units become available. A city councillor argues that the average wait for housing is therefore 8 years. A tenant advocate says the true average wait is much longer. Who is more likely to be right, and why? What would you need to know to compute the true average wait? What implicit assumption does the councillor's calculation make, and is it likely to hold?

Reflections

[Posted as a reflection exercise on the course's website for students to write about in a journal due at the end of semester.]

1. Think of something you believe to be true. It can be about people, about organizations, about how the world works that is not based on numerical data. What is the substance of the belief? What is it actually based on, if not data? How confident are you in it, and what would it take to change your mind?

2. Do you do anything in your daily life “suboptimally” from a personal standpoint, in order to benefit a group you belong to?
3. Think of a project—at work, in school, or elsewhere—where someone involved optimized their own piece of the work in a way that made the overall project harder or slower. Did they realize they were doing this? Did you say anything? Looking back, what might you have done differently to get a better outcome?
4. Ming argued in class that people are simply selfish and that voluntary compliance with the standing-only policy would never work. Do you agree? Can you think of examples from your own experience where people voluntarily accepted a personal cost for the benefit of a shared system, and examples where they did not? What seems to determine which way things go?
5. Recall a time you rushed hard to make it somewhere only to end up waiting in the end. Did your rushing help anyone, including yourself?
6. Camila resented being told to stand still when she could have walked. Name a small freedom of your own whose *universal* exercise would leave everyone, yourself included, worse off. Would you give it up voluntarily, or only if you were made to?
7. Ming struggled to accept a conclusion until she had seen the data, even when the logic was already airtight. Think of a time you demanded proof for something reason had settled, or trusted a pattern you never bothered to check. Which of those two mistakes do you make more often?
8. Little’s Law says the time you spend in any system is your inventory divided by the rate things clear. So if a line feels long, either there are many ahead of you or the line is barely moving. Think of a wait in your life that felt interminable. Was it long because of how many were ahead of you, or because of how slowly things were moving? Does knowing which one change how you’d feel about it?
9. One way to spend less time stuck in a system is to lower its inventory; another is to raise its throughput. Pick something in your life you’re currently “in the middle of”—a degree, a job search, a waitlist, a pile of work. Which lever is actually available to you, and which one are you tempted to reach for even though it isn’t?